CHAPTER 7

Intergenerational Altruism and the Effectiveness of Fiscal Policy—New Tests Based on Cohort Data

Andrew B. Abel and Laurence J. Kotlikoff

In recent years Barro's (1974) ingenious model of intergenerational altruism has taken its place among the major theories of consumption and saving. The model, which starts with the simple assumption that parents care about the welfare of their children, yields the remarkably strong conclusion that, apart from distorting marginal incentives, deficits and all other government redistributions between generations have no effect on the economy. The possibility that deficits, unfunded social security, and similar policies do not matter has received considerable attention.

Despite its policy importance, there have been few direct tests of the intergenerational altruism model. The main difficulty in directly testing the model at the microlevel is the relative lack of data detailing both the consumption and resources of altruistically linked households. In addition, it is difficult to determine from the data which households are altruistically linked to each other. Direct tests of the model with macrodata are also problematic because they require the aggregation of different clans (sets of altruistically linked households) each of which may have a different utility function.

In this essay we present a new direct test of the altruism model. The test

The earlier version of this paper (entitled "Does the Consumption of Different Age Groups Move Together? A New Nonparametric Test of Intergenerational Altruism," NBER Working Paper, no. 2490, 1988) appears to have been the first in a series of papers by economists, including Townsend (1989), Altonji, Hayashi and Kotlikoff (1992), Cochran (1991), and Mace (1991) to test the equality in Euler errors implied by risk sharing, whether the risk sharing is altruistically motivated or not. We thank Joseph Altonji, Robert Barsky, Gary Becker, Jinyong Cai, Gary Chamberlain, Jagadeesh Gokhale, Fumio Hayashi, Robert Lucas, Kevin Murphy, Sherwin Rosen, Robert Townsend, Lawrence Weiss, and seminar participants at Brown University, the University of Chicago, and at the National Bureau of Economic Research's summer workshop on financial markets for very helpful comments. Jinyong Cai and Jagadeesh Gokhale provided outstanding research assistance. This research was funded by the National Science Foundation. Our data are from the 1980 through 1985 Consumer Expenditure Surveys, which are publicly available through the Bureau of Labor Statistics.

is based on a property of the model that, as of the first draft of this essay, had not previously been exploited. This property is that the Euler errors (i.e., disturbances in the Euler equations) of altruistically linked members of clans are identical. Assuming utility is homothetic and time separable, this equality of Euler errors means that, controlling for clan preferences about the age distribution of consumption, the percentage changes over time in consumption of all clan members are equal. Intuitively, since consumption of each clan member is based on overall clan resources, and not on the distribution of resources over clan members, any shocks to the resources of specific clan members will be spread across all clan members. Under the homotheticity and time separability assumptions, spreading shocks over all clan members means changing the consumption of all members by the same percentage.

Ideally, one would test this proposition by simply comparing changes in the consumption of different clan members. Unfortunately, the requisite clan-specific data is not generally available (see Altonji, Hayashi, and Kotlikoff 1992 for an exception); indeed, it may be very difficult to determine who is and who is not a member of a particular altruistically linked clan. As indicated by Kotlikoff (1989) and Bernheim and Bagwell (1988), clans may be quite large because of current as well as potential intermarriage.

How can we use our Consumer Expenditure Survey data on household consumption to test intergenerational altruism when we cannot identify which households should be grouped together in clans? Our test is based on the implication of the altruism model that, after controlling for demographics, all clan members should change their consumption by the same percentage in any given period. If clans are large, and all clans have the same age structure, then the average Euler error in each age group of households in our sample should be the same, under intergenerational altruism. However, the assumption of identical age structures within each clan seems too strong. A weaker assumption is that the age structure of clans is independent of their Euler errors; that is, the fact that a clan accounts for a larger than average fraction of households in an age group does not help predict how its Euler error will differ, on average, from the average Euler error across clans. Even with this weaker assumption we can test the intergenerational altruism hypothesis by comparing the average Euler errors of different age groups.

Testing the altruism model by comparing average cohort percentage changes in consumption is particularly advantageous because it is non-parametric; in determining whether the average consumption of different age cohorts moves together we place no restrictions on preferences beyond the assumptions of homotheticity and time separability. In particular, each clan can have quite different preferences.

The new quarterly Consumer Expenditure Survey (CES), which, as of the time of this study, is available from the middle of 1980 through the middle of 1985, provides an excellent data set for determining whether the consumption of different age groups moves together. The CES records the consumption of each sample household for up to four quarters, and thus can be used to determine the average quarterly percentage change in consumption of households in a given age group.

The null hypothesis of our test is that, after controlling for demographics, cohort differences in the average percentage change in consumption (average Euler errors) are due simply to sampling and measurement error. Alternative hypotheses, suggested by the life-cycle model, are that, after controlling for demographics, (1) the percentage changes in the average consumption of any two cohorts are more highly correlated the closer in age are the two cohorts, (2) the variance in the percentage change in consumption is a monotone function of the age of the cohort, and (3) cohort differences in consumption changes depend on cohort differences in resource shocks, which may be proxied by cohort differences in income changes. While the data do not reject the altruism model against alternatives (1) and (2), they do reject the altruism model against alternative (3); that is, we find no age structure to the variance-covariance matrix of cohort Euler errors, but we do find that cohort differences in Euler errors are significantly correlated with cohort differences in income changes.

This essay proceeds in the next section by reviewing briefly the empirical literature bearing on the intergenerational altruism hypothesis. The third section presents the model of altruistically linked households and develops the proposition that Euler errors are equal for all clan members. The fourth section derives a statistical model to test this proposition. The fifth section discusses the statistical implications of the selfish life-cycle model, which represents at least one important alternative to the altruism model. The sixth section describes the data. The seventh section contains the empirical results. The last section summarizes the findings and concludes with a discussion of the implications of the findings for the effectiveness of fiscal policy.

Empirical Research Bearing on the Altruism Hypothesis

The largest body of empirical literature bearing on the altruism hypothesis relates the time series of aggregate consumption to the time series of social security wealth. Chief among these studies are those of Feldstein (1974), Darby (1979), and Leimer and Lesnoy (1980). Studies relating the consumption time series to other aspects of fiscal policy include Kormendi (1983), and Aschauer (1985). The results of this body of research can be summarized with

one word: ambiguous. Even were the results all in agreement, it would be difficult to know precisely what had been learned; as pointed out by Auerbach and Kotlikoff (1983) and Williamson and Jones (1983), if the life-cycle model is taken as the null hypothesis in these studies, the models are misspecified because of the inability to aggregate the behavior of different age groups. Auerbach and Kotlikoff (1983) show that the regression procedures would reject the life-cycle model even using data taken from a pure life-cycle economy. An alternative view of these regressions is that the Barro model is the null hypothesis. But in this case many of the regressions also seem to be misspecified both because of aggregation and because they ignore the government's intertemporal budget constraint.

Recent papers by Boskin and Kotlikoff (1986) and Boskin and Lau (1988) directly test the implication of the intergenerational altruism model that the age distribution of resources does not affect the age distribution of consumption. Both reject the proposition that aggregate consumption is invariant to the age distribution of resources. The findings of these two papers accord with the findings reported here that link cohort differences in Euler errors to cohort differences in income changes. Mace (1991) uses the same CES data that we use to test risk sharing. Her findings in support of risk sharing are strongly contradicted by the results reported here as well as by the findings in Altonji, Hayashi, and Kotlikoff 1989.

A different body of literature that is relevant to the altruism model as well as other neoclassical models is the Euler equation studies of Hall (1978), Flavin (1981), Hall and Mishkin (1982), Mankiw, Rothenberg, and Summers (1982), Lawrence (1983), Shapiro (1984), Altonji and Siow (1987), Zeldes (1989), and others. These studies test intertemporal expected utility maximization, specifically its implication that the Euler error is uncorrelated with previous information. A rejection of this null hypothesis would rule out the altruism model as well as other neoclassical consumption models. The time series tests of the Euler equation provide mixed results. In contrast, most microlevel studies appear to accept the Euler equation restriction for the majority of households. For example, both Zeldes (1989) and Lawrence (1983) use the limited consumption data in the Panel Study of Income Dynamics and reach the conclusion that the Euler equation holds for the great majority of households.

The microlevel studies that are closest to our own are Altonji, Hayashi, and Kotlikoff 1992, Townsend 1989, and Cochrane 1991. Altonji, Hayashi, and Kotlikoff use matched consumption and income data on the households of parents and their adult children; their findings strongly reject the altruism model's prediction that the distribution of consumption between parents and children is independent of the distribution of resources between the parent and children. In contrast, Townsend (1989) and Rosenzweig (1988) study the consumption and transfer behavior of households within Indian villages. Townsend reports that the consumption changes of households within each village are highly correlated and Rosenzweig reports substantial income-smoothing transfers; these findings are consistent either with altruism or risk sharing among villagers. Cochran tests for perfect insurance markets with microdata on consumption and certain types of income changes. Not surprisingly, he rejects the proposition that insurance markets are perfect, although he does report that certain income shocks are reasonably well insured.

Finally, there is a microliterature on transfers (see Cox 1987 and Kotlikoff 1988 for summaries) that appears, on balance, to reject the altruism model. Cox (1987), for example, finds that transfers rise with the level of recipients' incomes, and Menchik (1984) reports that, far from being equalizing, most bequests are divided perfectly evenly between children.

The Equal Euler Error Proposition

In this section we model the consumption decisions of households and show that all households within an altruistically linked clan will have the same Euler errors in every period. To model the consumption decisions we must introduce the utility function of the clan. The clan utility function depends on the consumption of present and future households within the clan. We begin by discussing the utility function of a household within a clan at a point in time. After discussing this intratemporal utility function, we combine the intratemporal utility functions of all households in the clan at all points of time.

The Intratemporal Household Utility Function

Let U_{ikt} denote the utility in period t of household k in clan i. The assumptions that the intertemporal utility function is homothetic and time-separable implies that the intratemporal household utility function is isoelastic; that is,

$$U_{i,k,t} = \sum_{a=0}^{D} P_{i,k,t,a}\,\theta_{i,k,t,a}\,\frac{c_{i,k,t,a}^{1-\gamma_i}}{(1-\gamma_i)}, \tag{1}$$

where $P_{i,k,t,a}$ is the number of members age a in household k, clan i at time t; D is the maximum age of life; $\theta_{i,k,t,a}$ is the weight household k in clan i places on the utility of members age a at time t; and $c_{i,k,t,a}$ is the consumption of the members of clan i who are in household k and are age a at time t. The utility function in equation 1 is written as a function of the consumption of each of the members of the household. For the isoelastic utility function

we can express household utility simply as a function of total household consumption at time t. Let $C_{i,k,t}$ denote household k's total consumption at time t, then:

$$C_{i,k,t} = \sum_{a=0}^{D} P_{i,k,t,a}\, c_{i,k,t,a}. \tag{2}$$

The optimal allocation of consumption across the individual members of the household is determined by maximizing the intratemporal utility function in equation 1 subject to equation 2. Performing this optimization and substituting the optimal values of $c_{i,k,t,a}$ into equation 1 yields

$$U_{i,k,t} = \phi_{i,k,t_i} \frac{C_{i,k,t}^{1-\gamma_i}}{(1-\gamma_i)} \tag{3a}$$

where

$$\phi_{i,k,t} = \left(\sum_{a=0}^{D} P_{i,k,t,a}\, \theta_{i,k,t,a}^{1/\gamma_i} \right)^{\gamma_i}. \tag{3b}$$

Equation 3a expresses the utility of the household as a function of total household consumption.

The Clan's Intertemporal Utility Function

The intertemporal utility function of a clan of altruistically linked households is obtained by summing the intratemporal household utility functions across all households in the clan and across the present and all future periods of time, taking account of time preference. Let α_i denote the time preference discount factor for all households in clan i, and let N_{is} denote the number of households in clan i at time s. At time t, the objective function of the clan is

$$V_{i,t} = E_t \sum_{s=t}^{\infty} \alpha_i^{s-t} \sum_{h=1}^{N_{i,s}} \phi_{i,h,s} \frac{C_{i,h,s}^{1-\gamma_i}}{1-\gamma_i}. \tag{4}$$

The clan maximizes equation 4 subject to:

$$W_{i,t+1} = (W_{i,t} + e_{i,t} - C_{i,t} - G_t)(1 + r_{i,t}), \tag{5}$$

where,

$$C_{i,t} = \sum_{h=1}^{N_{i,t}} C_{i,h,t},$$

is total clan i consumption at time t. The term $e_{i,t}$ stands for the possibly uncertain labor earnings of the clan at time t; $r_{i,t}$ is the possibly uncertain rate of return earned by clan i at time t on its portfolio of nonhuman wealth, and $W_{i,t}$ is clan i's net nonhuman wealth at time t. In addition to $e_{i,t}$ and $r_{i,t}$, $\phi_{i,h,s}$ for $s > t$ in equation 4 may be uncertain at time t due to life span uncertainty and uncertainty about clan fertility.

The term G_t in equation 5 stands for government consumption at time t. This is the only way in which fiscal policy enters the clan's budget constraint. Government consumption spending produces an income effect on private consumption spending, because the clan must finance this spending out of its (the economy's) collective output. But regardless of the size of government spending, the budget constraint is invariant to the choice by the government of which clan members will "pay" (mail the government its tax payments) for its spending. The reason is that the clan does its own redistribution across its members, so any redistribution by the government is automatically offset by clan redistribution. The fact that equation 5 does not include any information about the identity of taxpayers automatically implies that government redistribution in this model is ineffective.

The First-Order Conditions

Maximization of equation 4 subject to equation 5 implies the static first-order conditions

$$\phi_{i,k,t}\, C_{i,k,t}^{-\gamma_i} = \phi_{i,h,t}\, C_{i,h,t}^{-\gamma_i}, \tag{6}$$

and the intertemporal first-order conditions

$$E_t[\alpha_i \phi_{i,k,t+1}\, C_{i,k,t+1}^{-\gamma_i} (1 + r_{i,t+1})] = \phi_{i,k,t}\, C_{i,k,t}^{-\gamma_i}. \tag{7}$$

The static first-order conditions in equation 6 characterize the optimal allocation of consumption between households h and k within clan i. The intertemporal first-order condition in equation 7 holds for each household within clan i. Let $\epsilon_{i,k,t+1}$ denote the Euler error at time $t+1$ for household k in clan i. The Euler error is defined by

$$\alpha_i \phi_{i,k,t+1}\, C_{i,k,t+1}^{-\gamma_i} (1 + r_{i,t+1}) = E_t[\alpha_i \phi_{i,k,t+1}\, C_{i,k,t+1}^{-\gamma_i} (1 + r_{i,t+1})]\epsilon_{i,k,t+1}, \tag{8}$$

where $E_t\epsilon_{i,k,t+1} = 1$. Using the definition of the Euler error in equation 8 we can rewrite equation 7 as

$$\alpha_i \phi_{i,k,t+1} C_{i,k,t+1}^{-\gamma_i} (1 + r_{i,t+1}) = \phi_{i,k,t} C_{i,k,t}^{-\gamma_i} \epsilon_{i,k,t+1}. \tag{7'}$$

Equations 6 and 7′ together imply that the Euler errors of all households in clan i are identical. That is,

$$\epsilon_{i,k,t+1} = \epsilon_{i,h,t+1} \equiv \epsilon_{i,t+1}. \tag{9}$$

Note that in deriving this result we did not need to restrict different households in a given clan to have identical age compositions or to receive identical weights in the clan utility function.

Are Euler Errors Equalized under the Life-Cycle Model with Perfect Risk Sharing?

An alternative nonaltruistic model with the implication of equal Euler errors across households is the selfish life-cycle model under the assumption of perfect risk sharing and identical isoelastic preferences and identical time preference rates. To see this, note that the equilibrium of such a life-cycle economy with perfect risk sharing may be represented as the solution to a planning problem in which the planner maximizes a weighted sum of individual household expected intertemporal utilities subject to a collective budget constraint (see Townsend 1989). The division of total consumption at each point in time will depend on the household's weight; that is, equation 6 will hold. In addition, since there is a single budget constraint in this problem, the weighted marginal utility of each household's consumption will be equated to the shadow value of this budget constraint at each point in time. Hence, changes (over time) in the weighted value of each household marginal utility will equal the changes (over time) in these shadow prices; i.e., equation 7′ will hold. Since the shadow prices and their changes are not household specific, the weighted ratio of changes in marginal utilities of consumption will be the same for each household. Hence, equation 9, the equal Euler error proposition, will hold.

A Test of the Equal Euler Error Proposition Based on Cohort Data

In this section we develop a method of testing the equal Euler error proposition using cohort data. We start by taking logarithms of equation 7′, yielding

$$\log(C_{i,k,t+1}/C_{i,k,t}) = (1/\gamma_i)\log(\alpha_i \phi_{i,k,t+1}/\phi_{i,k,t})$$

$$- \log[\epsilon_{i,t+1}/(1 + r_{i,t+1})] \tag{10}$$

Consider all households in clan i whose heads are age a. Take the average of equation 10 over all such households. The resulting average of equation 10 is given by equation 11 where we define the averages of the left hand side and the two terms on the right hand side of equation 10 respectively by

$$Y_{i,t+1}^a = \psi_{i,t+1}^a + \mu_{i,t+1}. \tag{11}$$

Note that the term $\mu_{i,t+1}$ is not indexed by age since the Euler errors of each household in clan i are identical. Next average equation 11 over all clans. This produces equation 12 where $s_{i,t}^a$ is the fraction of age a households that belong to clan i at time t, and M is the total number of clans.

$$\sum_{i=1}^{M} s_{i,t}^a Y_{i,t}^a = \sum_{i=1}^{M} s_{i,t}^a \psi_{i,t}^a + \sum_{i=1}^{M} \mu_{i,t}/M + \sum_{i=1}^{M} s_{i,t}^a(\mu_{i,t} - \sum_{j=1}^{M} \mu_{j,t}/M) \tag{12}$$

In equation 12 the cohort average value of $\mu_{i,t}$ is written as the simple unweighted average of the Euler errors across all clans (the second term on the right hand side of the equation) plus the cohort average value (weighted by each clan's fraction of all cohort households) of the deviation of the clan's Euler error from the unweighted average Euler error over all clans. We assume that this third term on the right hand side, which is the population covariance between a clan's Euler error and its share of the population in the age group, is zero.

We can rewrite the remaining terms in equation 12 more compactly by letting $\bar{Y}_t^a$ denote the left hand side of equation 12, $\bar{\psi}_t^a$ denote the first term on the right hand side of equation 12, and $\bar{\mu}_t$ denote the second term on the right hand side of equation 12.

$$\bar{Y}_t^a = \bar{\psi}_t^a + \bar{\mu}_t \tag{12'}$$

Equation 12′ states that the cohort average value of the percentage change in consumption (more precisely, the log of the ratio of consumption at $t + 1$ to consumption at time t) equals a term, $\bar{\psi}_t^a$, which depends on age and time, plus a term $\bar{\mu}_t$, which is independent of age.

Because of sampling and measurement error, the true population mean, $\bar{Y}_t^a$, is not observable. Hence, in equation 13, we set the observed population-

weighted sample mean of the logarithm of the ratio of consumption at time $t + 1$ to consumption at time t, $\hat{Y}_t^a$, equal to the true population mean, $\bar{Y}_t^a$, plus a term, η_t^a, that reflects sampling and measurement error. Our null hypothesis is that $\eta_t^a = \omega_t^a / h_t^a$, where ω_t^a is an independently and normally distributed random variable with mean zero and variance σ^2, and h_t^a adjusts for the sampling error in our weighted estimate of $\bar{Y}_t^a$. Specifically, h_t^a equals $\Sigma_k w_{tk}^{a2} / (\Sigma_k w_{tk}^a)^2$, where w_{tk}^a is the CES population weight at time t for household k in cohort a. In equation 12′ the term $\bar{\psi}^a$, reflects the average growth in consumption due to demographic changes in household composition. Since we are dealing with data over only a five year interval, in equation 13 we drop the time subscript and treat $\bar{\psi}_t^a$ as a time-invariant, but age-specific constant.

$$\hat{Y}_t^a = \bar{\psi}^a + \bar{\mu}_t + \eta_t^a \tag{13}$$

Equation 13 forms the basis for our statistical test of the equality of average cohort percentage changes in consumption. Under the null hypothesis of equal Euler errors, ω_{at} is i.i.d. across ages a and time periods t with variance equal to σ^2.

If the null hypothesis fails to hold and the weighted average Euler errors differ across age cohorts, the error term η_t^a will capture not only measurement and sampling noise, but also each cohort's time t average Euler error after controlling for age and time effects. Our alternative hypothesis is, therefore, that the ω_{at}s are not simply i.i.d., but depend on age as specified below.

$$E(\omega_{it}\omega_{js}) = 0 \text{ if } s \neq t$$

$$E(\omega_{it}\omega_{jt}) = \rho^{|i-j|}\sigma^2\nu^{i+j} \tag{14}$$

According to equation 14 the variance of ω_{it} increases or decreases with age depending on whether ν exceeds or falls short of unity, and the correlation of ω_{it} and ω_{jt} for $i \neq j$ depends on the size of the age gap, $|j - i|$. For example, if ρ exceeds zero, equation 14 says that the correlation of ω_{it} and ω_{jt} for age groups i and j is larger the closer in age are the age groups i and j. The case in which $\rho = 0$ and $\nu = 1$ corresponds to the null hypothesis. Values of ρ and ν as well as the age and time effects in equation 13 are estimated by maximum likelihood. The Appendix presents the likelihood function and derives the estimators.

Another testable implication of equation 13 is that η_t^a is uncorrelated with changes in cohort a's resources, which may be proxied by changes in its income. To see this note that in equation 13 the term $\bar{\mu}_t$, which equals the common (across cohorts) average Euler error, fully controls for resource changes under the altruism model. Hence, if one adds cohort a's income change to the implicit regression model in equation 13, the coefficient on the cohort's income change should be zero. Another way of saying this is that differences across cohorts with respect to consumption changes should depend only on differences in their demographics (the $\bar{\psi}^a$ terms) and not on the distribution across cohorts of income changes. In addition to testing whether there is a significant age pattern to the variance-covariance matrix of the η_t^as, we add the cohort's income change to equation 13 and estimate the model by ordinary least squares. This procedure is, in differences, the fixed effects test of altruism developed in Altonji, Hayashi, and Kotlikoff 1989.

Consumption Behavior of Selfish Life-Cycle Households: An Alternative Hypothesis

This section motivates the assumption of an age-dependent variance-covariance matrix of the η_t^as under the alternative life-cycle model. The null hypothesis of operative altruistic linkages is that the Euler error is identical across households within a clan; hence, within a clan the Euler error is independent of the age of the household head. In contrast, under the life-cycle model Euler errors of different households within a clan bear no special relation to one another, but we would expect that the variance-covariance matrix of Euler errors across unrelated as well as related households would depend on the households' ages. This section illustrates, with two different preference structures and types of uncertainty, why ρ is likely to differ from zero and ν is likely to differ from unity if the life-cycle model holds.

Example 1: Logarithmic Utility

The first example is based on Samuelson 1969 and assumes only uncertainty with respect to the rate of return. Let $c_{t,a}$ be the consumption at time t of a household whose head is age a; $w_{t,a}$ is the wealth at time t of a household whose head is age a. The decision problem at time t, when the household is age a, is to maximize

$$E_t\left[\sum_{j=0}^{D-a} \beta^j u(c_{t+j,a+j})\right], \tag{15}$$

and

$$w_{t+1,a+1} = (w_{t,a} - c_{a,t})R_{t+1,a+1}, \tag{16}$$

where D is the age at which the household head dies and $R_{t+1,a+1}$ is the gross rate of return on the portfolio from period t to period $t + 1$. Now suppose that the utility function is logarithmic, $u(c) = \log c$. This optimization problem

can be solved by stochastic dynamic programming (Samuelson 1969) to obtain

$$c_{t,D-j} = g_j w_{t,D-j}, \tag{17a}$$

where

$$g_j = \left(\sum_{k=0}^{j} \beta^k \right)^{-1}. \tag{17b}$$

This solution holds regardless of the temporal dependence of the process generating returns $R_{t+1,a+1}$.

Now consider the growth rate in consumption from period t to period $t + 1$. It follows directly from equation 17a that

$$c_{t+1,a+1}/c_{t,a} = (g_{j-1}/g_j)(w_{t+1,a+1}/w_{t,a}). \tag{18}$$

The ratio of wealth in successive periods can be rewritten using equations 16 and 17a as

$$w_{t+1,a+1}/w_{t,a} = (1 - g_j)R_{t+1,a+1}. \tag{19}$$

Substituting equation 19 into equation 18 yields

$$c_{t+1,a+1}/c_{t,a} = (g_{j-1}/g_j)(1 - g_j)R_{t+1,a+1}. \tag{20}$$

Finally, we can use the expression for g_j in equation 17b to simplify the expression for the growth rate of consumption in equation 20 to obtain

$$c_{t+1,a+1}/c_{t,a} = \beta R_{t+1,a+1}. \tag{21}$$

The growth rate of consumption is proportional to the realized gross rate of return from period t to period $t + 1$. If the gross rate of return on a household's portfolio is independent of the age of the household head then $R_{t+1,a+1} = R_{t+1}$ for all a. In this case, all households will have the same Euler error regardless of the age of the household head.

Taken at face value, this example suggests that the Euler error is independent of the age of the household head. However, this conclusion depends on the assumption that $R_{t+1,a+1} = R_{t+1}$ for all a. This assumption would be warranted if all households held the same portfolios (up to a scale factor) regardless of age. However, as documented in King and Leape 1984 the composition of actual U.S. household portfolios is significantly different depending on the age of the household head. Young and middle-age households tend to hold a large portion of their wealth in the form of negative holdings of fixed income securities (mortgages), significant holdings of housing, and small holdings of stocks and bonds. Older households, in contrast, hold much more of their wealth in the form of home equity (i.e., their outstanding mortgages are much smaller) and in stocks and bonds.

Because the allocation of portfolios varies systematically with age, and because the rates of return on different assets reflect different stochastic processes, we might expect the conditional variances of the portfolio rates of return, $R_{t+1,a+1}$, to vary systematically with age. Furthermore, because the composition of portfolios is more similar for similar aged households than for households of very different ages, we might expect the Euler errors to be more highly correlated for households of similar age than for households of very different ages.

Example 2: Quadratic Utility

In order to focus on the role of human wealth, as distinct from nonhuman wealth, we change the framework slightly. Now we suppose that utility is quadratic and that the rate of return on nonhuman wealth is constant. In addition, assume that the gross rate of return on nonhuman wealth, R, is equal to β^{-1} (the reciprocal of the time preference discount factor). The only uncertainty that the household faces is in labor income $y_{t,a}$. In this case, it is straightforward to apply the certainty equivalence principle to obtain:

$$E_t(c_{t+j,a+j}) = c_{t,a} \text{ for } j = 0, 1, 2, \ldots, D - a. \tag{22}$$

The lifetime budget constraint of the household implies that the present value of revisions in future labor income, $E_{t+1}(y_{t+1+j,a+1+j}) - E_t(y_{t+1+j,a+1+j})$, is equal to the present value of revisions in future consumption, $E_{t+1}(c_{t+1+j,a+1+j}) - E_t(c_{t+1+j,a+1+j})$. Such that

$$\sum_{j=0}^{D-a-1} R^{-j}[E_{t+1}(y_{t+1+j,a+1+j}) - E_t(y_{t+1+j,a+1+j})]$$

$$= \sum_{j=0}^{D-a-1} R^{-j}[E_{t+1}(c_{t+1+j,a+1+j}) - E_t(c_{t+1+j,a+1+j})]. \tag{23}$$

Equation 22 implies that the revisions in expectations of future consumption are equal to the change in consumption between period t and period $t + 1$, hence:

$$E_{t+1}(c_{t+1+j,a+1+j}) - E_t(c_{t+1+j,a+1+j})$$

$$= c_{t+1,a+1} - c_{t,a} \equiv \Delta c_{t+1,a+1}. \tag{24}$$

To calculate the revision in expectations of future labor income, we first specify the moving average representation of the process for labor income as

$$y_{t,a} = \sum_{k=0}^{a} \xi_k e_{t-k,a-k} + \ddot{y}_{t,a}, \tag{25}$$

where $E_{t-1}\{e_{t,a}\} = 0$. With this time series process for labor income, the revisions in expected future (between time t and time $t + 1$) labor income at time $t + 1$ are

$$E_{t+1}(y_{t+1+j,a+1+j}) - E_t(y_{t+1+j,a+1+j}) = \xi_j e_{t+1,a+1}. \tag{26}$$

Substituting the revisions in future consumption equation 24 and the revisions in future labor income equation 26 into equation 23 we obtain

$$\Delta c_{t+1,a+1} = \Gamma_a e_{t+1,a+1}, \tag{27a}$$

where

$$\Gamma_a \equiv \left(\sum_{j=0}^{D-a-1} R^{-j} \xi_j \right) / \left(\sum_{j=0}^{D-a-1} R^{-j} \right). \tag{27b}$$

The variance of the unforecastable change in consumption, which in this example is equal to the variance of the actual change in consumption, is

$$\text{var}(\Delta c_{t+1,a+1}) = \Gamma_a^2 \text{ var}(e_{t+1,a+1}). \tag{28}$$

Note that even if the variance of the innovation to the labor income process, $\text{var}(e_{t+1,a+1})$, is independent of age, the variance of the unforecastable change in consumption is age-dependent because Γ_a is, in general, age-dependent. For instance, if the labor income process is i.i.d., then $\xi_0 = 1$ and $\xi_j = 0$ for all nonzero j. In this case,

$$\Gamma_a = \left(\sum_{j=0}^{D-a-1} R^{-j} \right)^{-1},$$

which is an increasing function of age. Alternatively, if the labor income process is a first-order auto-regressive progress with AR coefficient ρ, then $\xi_j = \rho^j$. In this case, $\Gamma_a = [(1 - R^{-1})/(1 - \rho/R)][1 - (\rho/R)^{D-a})/(1 - (1/R)^{D-a})]$. If the process is stationary and ρ is nonnegative, then Γ_a is increasing with age. For a random walk, $\rho = 1$ and $\Gamma_a = 1$ independent of age. If ρ is greater than 1, then Γ_a is decreasing with age.

Equation 28 gives the variance of the Euler error expressed in terms of the change in the level of consumption rather than the change in the logarithm of consumption. In this model, the age-consumption profile will, on average, be flat. Hence, on average, the variance of the percentage change in consumption will vary with age if the variance of the absolute change in consumption varies by age.

The Data

The ongoing Consumer Expenditure Survey (CES), which began in the first quarter of 1980, interviews approximately 4,500 households in each quarter. Most households are interviewed four times in the CES. The four interviews always ask a common set of questions about consumption, but some questions are asked only in the first and fourth interviews, and others are asked only in the fourth interview. Some households are interviewed fewer than four times because they drop out of the sample. Others are interviewed fewer than four times because of the sample design; in an effort to maintain in each quarter the same fraction of households responding to a first, second, third, and fourth interview, the CES administers the second, third, or fourth interviews to some households as their initial interview. If the household's initial interview is a second interview, the household will be interviewed two more times. If a household's initial interview is a third interview, the household will be interviewed once more. And if the household's initial interview is a fourth interview, the household will not be reinterviewed.

The approximately 4,500 interviews in each quarter are spread over each month of the quarter. In the interviews households are asked about their consumption expenditures in the previous three months. Hence, a household interviewed in January 1981 reports consumption expenditures for October, November, and December 1980, while a household interviewed in March 1981 reports consumption expenditures for December 1980 and January and February 1981. Unfortunately, for most expenditure items, households only report total expenditures in the previous three full months and do not provide a month-to-month breakdown of those expenditures. As a consequence, the data for a household interviewed, say, in January cannot readily be combined with data from a household interviewed in February since the two quarterly observations cover overlapping, rather than identical quarters. In effect, each

wave of the Consumer Expenditure Survey provides three overlapping sets of observations on quarterly consumption. In our analysis we treat each of the three quarterly data sets separately and refer to them as "quarterly sample" 1, 2, and 3.[1] For purposes of analyzing the quarterly data we considered fifty-eight age cohorts corresponding to ages 23 through 80.

Given the lumpiness of some nondurable consumption expenditures, such as vacation trips, it is useful to test the equal Euler error proposition with semiannual as well as quarterly data. For those households who were interviewed four times, the four quarterly observations can be combined to form observations on semiannual consumption. There are six possible semiannual data sets. For example, households interviewed in January, April, July, and October in year t provide an observation on the ratio of consumption over the period April–September in year t to consumption over the period October in year $t - 1$–March in year t. Households interviewed in July and October of year t and January and April of year $t + 1$ provide an observation on the ratio of consumption over the period October in year t–March in year $t + 1$ to consumption over the period April–September in year t. These types of observations produce a single data set of semiannual changes in consumption. One can also form a data set using households interviewed for the first of four times in April and other households interviewed for the first of four times in October. Hence, the April-July-October-January sequence provides two semiannual data sets. The May-August-November-February sequence provides another two semiannual data sets; and the June-September-December-March sequence provides the final two semiannual data sets.

Because of the smaller number of households who completed all four surveys, we constructed three-year age cohorts; i.e., we combined ages 23, 24, and 25 into one age group, ages 26, 27, and 28 into another age group, etc., up to the age group covering ages 77, 78, and 79. This difference in the definition of an age cohort should be kept in mind when comparing the quarterly and semiannual results presented in the next section; because of the difference in definitions, one would expect the estimated values of ρ and ν based on the semiannual data to be roughly the cube of their respective values based on the quarterly data.

1. Quarterly sample 1 corresponds to households interviewed in April, July, October, and January. Quarterly sample 2 corresponds to households interviewed in May, August, November, and February. Quarterly sample 3 corresponds to households interviewed in June, September, December, and March. In constructing the data for quarterly sample 1, as an example, we form ratios of (*a*) the July reported quarterly consumption to the April reported quarterly consumption, (*b*) the October reported quarterly consumption to the July reported quarterly consumption, (*c*) the January reported quarterly consumption to the October reported quarterly consumption, and (*d*) the April reported quarterly consumption to the January reported quarterly consumption. In forming the average logarithm of the ratio of consumption say in January 1983 to consumption in October 1982, all households who were surveyed in both October 1982 and January 1983 were included.

The definition of aggregate consumption used in this study is total consumption expenditures excluding expenditures on housing, insurance, and consumer durables. We exclude housing both because adjustments to housing consumption are infrequent and because it is very difficult to impute quarterly or semiannual rent accurately for homeowners. Insurance expenditures were excluded because such expenditures represent risk pooling as opposed to consumption per se. In addition, the data records both negative and positive amounts of insurance expenditures, where a negative amount corresponds to a claim payment. Expenditures on durables should clearly be excluded from the definition of consumption. In contrast, imputed rent should be included; unfortunately, data on the stocks of durables are not sufficient for that purpose.

The CES provides population weights in each quarter for each household interviewed. These weights depend on the age of the household head as well as other economic and demographic characteristics. We use the time $t + 1$ sample weights in determining the cohort-specific weighted average value of the logarithm of the ratio of consumption at time $t + 1$ to consumption at time t; that is, we construct a weighted value of $\bar{Y}_{at}$.

Households that reported less than $150 of quarterly expenditure on food were excluded from the sample. This is the only form of sample selection in our analysis. Some preliminary analysis indicated that including households with very small quarterly food expenditure would not materially alter the results.

Empirical Findings

Changes in the Age-Consumption Profile over the Sample Period

As a prelude to examining estimates of ρ and ν, figure 1 illustrates how the age-consumption profile changed over the period 1980 through 1984. Ignoring demographic change, the proposition that each cohort's consumption should change, on average, by the same percentage, implies a time-invariant age-consumption profile. In forming figure 1 we calculated the annual weighted average of quarterly consumption (measured in 1985 dollars) at each individual age for households interviewed in April, August, and December of each of the five years. We combined these weighted averages within each calendar year to produce annual values of average consumption by age of the household. Next we divided annual consumption in year t at each age by the average consumption of 45-year-old households in year t. Finally, we smoothed these relative consumption values for each year by regressing them against an intercept and a fourth order polynomial in age. In these regressions the R^2 values each exceed 0.9. Figure 1 plots the resulting five smoothed polynomials of consumption at a particular age relative to consumption at age 45.

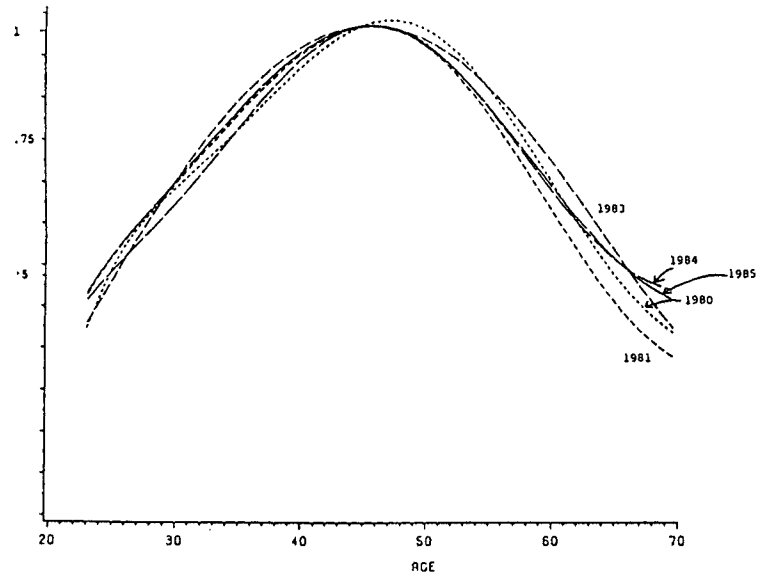

Fig. 1. Age consumption profile, 1980–85

The curve with the most dashes corresponds to 1980, the curve with the second most dashes corresponds to 1981, etc. The curves in the figure suggest that the age-consumption profile flattened out in 1982 and 1984. Compared with 1981, for example, the 1984 relative consumption of 60 year olds is over 10 percent larger. The $F(20, 264)$ value for the test that the five polynomials are the same is 17.94, greatly in excess of the 5 percent critical value of 1.66. To the extent that the changes in the shape of the relative age-consumption profile do not appear to be due to changes in demographics, it provides some evidence against the intergenerational altruism model; however, unlike the next set of findings, these profiles consider levels, not changes in consumption, and, as such, do not control as well for the composition of the sample; that is, the levels of consumption of the elderly in 1983 and 1984 may reflect samples whose older households happened to belong to clans with greater total resources.

Testing for Age-Dependence in Cohort Euler Errors

Table 1 also provides some preliminary data anlysis of the null hypothesis. This table compares quarterly changes in consumption of different age groups for quarterly sample 1. The corresponding tables for quarterly samples 2 and 3 are quite similar. For purposes of table 1 we consider five broad age catego-

TABLE 1. Quarterly Estimates of Average Adjusted Euler Errors Quarterly Sample I

Age Group	1980.3	1980.4	1981.1	1981.2
23–29	−.033 (0.204)	.055 (.178)	.074 (.178)	−.105 (.181)
30–39	−.101 (0.210)	.057 (.181)	.052 (.178)	−.087 (.174)
40–49	−.084 (0.198)	.100 (.182)	−.004 (.175)	−.082 (.192)
50–59	−.154 (0.179)	.041 (.175)	−.002 (.181)	−.027 (.183)
60–69	−.130 (0.193)	.015 (.220)	.024 (.169)	.006 (.207)
70+	−.109 (0.185)	.010 (.202)	.099 (.170)	−.020 (.202)
	1981.3	1981.4	1982.1	1982.2
23–29	−.050 (.187)	−.007 (.183)	−.115 (.182)	.031 (.175)
30–39	.056 (.181)	.017 (.189)	−.061 (.185)	−.055 (.193)
40–49	−.060 (.199)	.031 (.199)	−.076 (.179)	−.024 (.201)
50–59	−.074 (.180)	.004 (.181)	−.020 (.180)	−.032 (.179)
60–69	−.038 (.204)	−.071 (.223)	−.023 (.173)	−.094 (.199)
70+	−.154 (.198)	.105 (.236)	−.123 (.174)	.140 (.196)
	1982.3	1982.4	1983.1	1983.2
23–29	−.006 (.182)	.020 (.184)	.029 (.179)	−.048 (.177)
30–39	−.014 (.201)	−.010 (.197)	.111 (.182)	−.052 (.178)
40–49	−.008 (.200)	.008 (.195)	.088 (.177)	−.038 (.181)
50–59	−.062 (.180)	−.016 (.188)	.052 (.174)	−.020 (.209)
60–69	−.035 (.193)	−.077 (.202)	.049 (.175)	−.101 (.200)
70+	−.108 (.209)	.009 (.219)	.112 (.179)	−.110 (.222)
	1983.3	1983.4	1984.1	1984.2
23–29	.032 (.184)	−.003 (.175)	.015 (.176)	−.076 (.205)
30–39	−.031 (.181)	.011 (.230)	.008 (.187)	−.061 (.183)
40–49	−.004 (.183)	.065 (.180)	−.011 (.176)	−.060 (.181)
50–59	.042 (.207)	.012 (.183)	.011 (.174)	−.027 (.186)
60–69	−.020 (.221)	.012 (.217)	−.064 (.173)	−.003 (.181)
70+	−.020 (.212)	.019 (.196)	−.012 (.178)	.009 (.190)
	1984.3	1984.4	1985.1	
23–29	−.021 (.213)	.046 (.207)	−.011 (.188)	
30–39	.020 (.185)	.027 (.185)	.034 (.185)	
40–49	.025 (.183)	.034 (.185)	.032 (.186)	
50–59	.007 (.191)	−.034 (.194)	−.065 (.181)	
60–69	−.011 (.178)	−.046 (.180)	.014 (.176)	
70+	−.015 (.186)	−.029 (.190)	.016 (.176)	

Note: Standard deviations are in parentheses.

ries: 23–29, 30–39, 40–49, 50–59, 60–69, and 70+. For each of these six age groups we report quarterly values of 100 times the deviation of the weighted average $\hat{\bar{Y}}_t^a$ from the mean value, $\Sigma_{t=1}^{19} \hat{\bar{Y}}_t^a$, taken over the nineteen quarters in our sample. According to equation 12′ and ignoring measurement and sampling error, these deviations, which we refer to as average adjusted Euler errors, should be identical for each of the six age groups. In addition to presenting these deviations, table 1 reports the standard deviations of these deviations. These standard deviations are based on information on the variation across households in the percentage change in consumption within age-time cells.

Table 1 indicates that these average adjusted Euler errors are typically different across the six age categories. However, these differences could well be due to within cell variation. The standard errors of the entries in table 1 are quite large.

Another informal way to assess the data is to regress $\hat{\bar{Y}}_t^a$ on a set of age group dummies and time dummies, either quarterly or semiannual.[2] The results from this regression can be compared with the results from regressing the same dependent variable on age dummies and the interaction of each of the time dummies with each of the age dummies.[3] According to equation 13, given a particular time period t, the age-time interactions should have identical coefficients. For purposes of this regression using quarterly data, we constructed six age dummies corresponding to the six age groups of table 1. The F values for quarterly samples 1, 2, and 3 are 1.470, 1.237, 0.746, respectively. Since the $F(90, 987)$, 5 percent critical value is 1.27, the age-time interactions are significantly different in only one of the three quarterly samples. The $F(25, 103)$ values in the corresponding regressions for the six semiannual samples are 0.912, 2.414, 2.538, 1.485, 1.768, and 0.612. The 5 percent critical value in this case is 1.61. Hence, age-time interactions are significantly different in three of the six semiannual samples.

Table 2 presents our maximum likelihood estimates for ρ and ν for the three quarterly data sets based on individual age cohorts from age 23 through age 80. None of the reported estimates of these parameters is significantly different from the values predicted by the null hypothesis of intergenerational altruism. Indeed, in the case of ν, two of the three estimates are equal to 1, to three decimal places, and the third value of 1.002 implies that the variance of

2. In this regression time dummies, age dummies, and the residuals account, respectively, for 81.3 percent, 4.2 percent, and 14.5 percent of the variance of the dependent variable.

3. To illustrate these two regressions, consider the case of two age groups, ages 1 and 2, and two time periods, time 1 and 2. Then the initial regression is:

$\hat{\bar{Y}}_t^a = \delta_1 A_1 + \delta_2 A_2 + \tau T_2$, where the δs and τs are coefficients, and A_1, A_2 and T_2 are dummies for age group 1, age group 2, and time period 2, respectively. The alternative model is: $\hat{\bar{Y}}_t^a = \delta_1 A_1 + \delta_2 A_2 + \lambda A_1 T_2 + l A_2 T_2$, and the test is whether $\lambda = l = \tau$.

TABLE 2. Maximum Likelihood Estimates and χ^2 Values: Quarterly Samples

	Consumption Measured per Household						
	Unconstrained		$\rho = 0$		$\nu = 1$		$\rho = 0$ $\nu = 1$
Sample	ρ	ν	ν	χ^2	ρ	χ^2	χ^2
Sample 1	0.020	1.002	1.002	0.958	0.003	1.307	2.345
Sample 2	−0.024	1.000	1.000	0.611	−0.024	0.000	0.611
Sample 3	0.031	1.000	1.000	1.066	0.031	0.000	1.066

Note: Five percent critical values for χ^2 are 5.991 for two restrictions and 3.841 for one restriction.

ω_{it} for 70 year olds is only about 15 percent larger than the corresponding variance for 20 year olds. One of the three point estimates for ρ is negative; a negative value of ρ, even were it significant, seems highly unlikely from the perspective of the life-cycle model. The other two nonnegative values of ρ suggest a very small correlation between the consumption of adjacent age groups. Even if these estimates were significant, their values seem quite small.

The likelihood functions associated with table 2 are rather sharply peaked; hence, one can reject values of ρ and ν that are substantially different (in an economic sense) from the maximum likelihood estimates. Table 3 presents the range of values of ρ and ν that fall within 95 percent chi-squared confidence intervals around the maximum likelihood estimates.[4] According to the table, even if one takes the largest values of ρ and ν that cannot be rejected by the data, the resulting estimates provide no strong evidence of substantial departure from the null hypothesis of intergenerational altruism.

Since many of the consumption expenditures included in our definition of nondurable expenditure may not be made each quarter, the results in table 2 may, in part, reflect the lumpiness of nondurable expenditures; that is, the variance in consumption changes due to the lumpiness of expenditures may dominate the results. Hence, it may be useful to repeat the analysis using simply food expenditure, which is much less lumpy than, for example, clothing expenditure or vacation trips. The results based on quarterly food expenditures are quite similar to the results based on total nondurable, non-housing consumption expenditure. The point estimates in the three samples of ρ are −0.059, −0.068, and 0.002. The point estimates of ν in the three samples are 1.001, 0.998, and 1.003. The estimates of ρ and ν are not jointly significantly different from 0 and 1, respectively; the respective x^2 values for

4. These bounds were constructed by holding one of these parameters fixed at its maximum likelihood value and varying the other parameter until the resulting likelihood was significantly (at the 5 percent level) different from the maximum likelihood.

TABLE 3. 95 Percent Confidence Intervals for ρ and ν: Quarterly Estimates

	Consumption Measured per Household					
	Maximum Likelihood		ν Range		ρ Range	
Sample	ρ	ν	High	Low	High	Low
Sample 1	0.020	1.002	1.005	0.999	.095	−0.030
Sample 2	−0.024	1.000	1.003	0.997	.020	−0.099
Sample 3	0.031	1.000	1.003	0.997	.106	−0.044

the three samples for the joint test that $\rho = 0$ and $\nu = 1$ are 4.087, 4.991, and 4.757—all of which lie below the 5 percent critical value of 5.991.

Another way to consider the lumpiness of expenditures is to repeat the analysis with semiannual data. Table 4 presents the results based on the six semiannual consumption data sets, which, as mentioned, combine three ages into a single-age cohort. Once again, none of the estimates of ρ and ν are separately or jointly significantly different from the null hypothesis values of $\rho = 0$ and $\nu = 1$. Three of the six point estimates of ν lie above 1 and three lie below 1. Three of the six point estimates of ρ are positive and three are negative. Hence, like the quarterly estimates, there is no suggestion in the data that the null hypothesis is strongly disfavored. Unlike the quarterly results, however, several of the estimates of ν are economically more important. For example, the estimate for ν in the sixth sample of 1.022 implies that the variance of ω_{it} for very old households is over 1.7 times the variance for very young households. In addition, for each of the six samples the confidence intervals around ν include economically significant as well as economically insignificant values. Thus the semiannual results do not provide as

TABLE 4. Maximum Likelihood Estimates and χ^2 Values: Semiannual Samples

	Consumption Measured per Household						
	Unconstrained		ρ = 0		ν = 1		ρ = 0 ν = 1
Consumption Category	ρ	ν	ν	χ^2	ρ	χ^2	χ^2
Sample 1	−0.133	1.007	1.009	2.531	−0.136	0.409	3.101
Sample 2	−0.142	1.008	1.009	2.945	−0.142	0.478	3.490
Sample 3	0.087	0.990	0.991	0.024	0.080	0.706	1.585
Sample 4	0.055	0.991	0.992	0.467	0.049	0.555	0.918
Sample 5	−0.072	0.984	0.984	0.719	−0.069	2.377	3.054
Sample 6	0.050	1.022	1.022	0.353	0.051	3.679	4.048

Note: Five percent critical values for χ^2 are 5.991 for two restrictions and 3.841 for one restriction.

strong evidence against the life-cycle model as do the quarterly results. It may be that quarterly changes and even semiannual changes in consumption reflect quite lumpy expenditures and that testing the equal Euler error proposition on annual or even biannual would be more appropriate. Unfortunately, appropriate data for such an analysis do not currently exist.

One might question whether we have properly controlled for demographic change in treating $\bar{\psi}^a$ as an age-specific, time-invariant constant. One way to consider whether the results are sensitive to treatment of demographics is to reestimate the model defining household consumption as household consumption per household member or per adult equivalent in the household; in forming adult equivalents we treat each child under age 18 as equal to .5 adults. We tried each of these alternative definitions of household consumption. The quarterly results are essentially the same as those in table 2.[5] The semiannual results are only slightly different from those in table 4; when consumption is measured either as household consumption per member or per equivalent adult, the null hypothesis is rejected in only two of the six semiannual samples.[6]

Are Cohort Euler Errors Correlated with Cohort Changes in Income

One question that may be raised with respect to the empirical findings in the above subsection is the power of the tests against the life-cycle alternative. While we have argued that the variance-covariance matrix of cohort Euler errors is likely to be age-dependent under the life-cycle model, there are some combinations of life-cycle preference structures and distributions of cohort-specific resource shocks that also satisfy the null hypothesis of $\nu = 1$ and

5. With household consumption defined as consumption per person, the point estimates for ρ for samples 1, 2, and 3 are 0.034, −0.043, and 0.044, respectively. For ν the corresponding point estimates are 1.001, 1.000, and 1.000. With household consumption defined as consumption per equivalent adult, the three point estimates for ρ are 0.032, −0.035, and 0.039, while the three point estimates for ν are 1.001, 1.000, and 1.000. The χ^2 values for testing the null hypothesis that $\rho = 0$ and $\nu = 1$ are 2.094, 1.952, and 0.395 for the three quarterly samples when consumption is measured per person, and 2.240, 1.352, and 1.667 for the three quarterly samples when consumption is measured per equivalent adult.

6. With household consumption defined as consumption per person the point estimates for ρ for semiannual samples 1, 2, 3, 4, 5, and 6 are −0.1260, 0.0450, −0.0840, −0.1360, 0.0570, and −0.0060, respectively. For ν the corresponding point estimates are 1.0010, 0.9710, 0.9700, 1.0140, 0.9970, and 1.023. With household consumption defined as consumption per equivalent adult, the six point estimates for ρ are −0.1270, 0.0770, −0.0900, −0.1310, 0.0480, and .0100, while the six point estimates for ν are 1.0090, 0.9730, 0.9740, 1.015, 0.9940, and 1.0210. The six respective semiannual χ^2 values for testing the null hypothesis that $\rho = 0$ and $\nu = 1$ are 3.0763, 7.3924, 8.710, 3.807, 0.493, and 4.220 when consumption is measured per person, and 3.0441, 6.699, 7.284, 3.735, 0.533, and 3.581 when consumption is measured per equivalent adult.

TABLE 5. Cohort Average Income Change Coefficients

Sample	Regressions with Time Dummies		Regressions without Time Dummies	
	Coefficient	*t*-statistic	Coefficient	*t*-statistic
Quarterly sample 1	.035	1.52	.011	.38
Quarterly sample 2	.068	3.58	.027	.85
Quarterly sample 3	.066	2.70	.034	1.19
Pooled sample	.052	3.57	.026	1.53

$\rho = 0$. A potentially more powerful method for testing altruism against the life-cycle alternative is to ask whether, in addition to age dummies and time dummies (that control for common resource shocks), cohort-specific income changes enter significantly in a regression explaining cohort average percentage changes in consumption. Because the CES asked about income only in its first and fourth interviews, we only have two observations on income. Hence, we considered only households with four complete surveys and formed, for each of the three-year age groups, the average percentage change in income and consumption between the first and fourth interviews. We did this for each of the three quarterly samples.

Table 5 reports the income change coefficients for two different sets of regressions. In the first set we control for common resource shocks by including time dummies. In the second set we exclude the time dummies. Table 5 contains the results from the regressions for each of the quarterly samples, plus the results from pooling the three quarterly samples. In the case of the pooled regression with time dummies, we include different time dummies for each quarterly sample since, as described above, the quarterly samples cover somewhat different time intervals.

Consider the regressions including the time dummies. In two of three quarterly samples and in the case of the pooled regression, the coefficients on cohort income changes are statistically significant. In addition, the point estimates are economically large in the sense that they are even larger than the point estimates one obtains if one does not control for common resource shocks through the time dummies. These results constitute evidence against the altruism model and evidence for the life-cycle model.[7]

7. The simple Keynesian model, in which consumption depends only on contemporaneous income, also is supported by these findings. But there appears to be more support in the results for the life-cycle model than the Keynesian model. The reason is that the time dummies in the first set of regressions in table 5 are highly significant. Ignoring measurement error, the Keynesian model (at least the simple version of it) predicts that only changes in current income explain changes in consumption. In the Keynesian model, therefore, the time dummies, which control for common resource shocks, should not enter the regression explaining changes in cohort consumption. In contrast, the life-cycle model would suggest that the time dummies as well as the income change variable would be significant. The reason is that the income change variable does not, under the life-cycle hypothesis, control perfectly for the cohort's change in resources; if the life-cycle model holds and the Euler errors of different age groups are correlated, then the time dummies should pick up that component of the common shock to the different cohorts' Euler errors that are not captured by the change in income. This point was first made in Altonji, Hayashi, and Kotlikoff 1989.

The Implications of the Findings for the Effectiveness of Fiscal Policy

In this essay we have used cohort data to test the standard model of operative altruistic linkages. The altruism model in contrast to the standard (no-risk-sharing) life-cycle model suggests that (1) after controlling for demographics, consumption changes for different cohorts should, on average, be identical and (2) that, after controlling for demographics, differences across cohorts in their changes in consumption should not be correlated with their income changes. The life-cycle model plus the assumption of perfect risk sharing also delivers these two testable propositions. We tested the first implication by asking whether, after controlling for demographics and common resource shocks, the variance-covariance matrix of consumption changes is age-dependent; we tested the second implication by determining whether differences across cohorts in consumption changes depend on differences across cohorts in their income changes.

We were not able to reject the first implication of the altruism model, but we were able to reject the second. Hence, on balance, our results reject the altruism model and, by extension, the life-cycle model with risk sharing in favor of the standard (no-risk-sharing) life-cycle alternative. An important attribute of these results is that they are nonparametric in nature; specifically, in comparing the average change in consumption across age groups, we place no restrictions on preferences beyond the assumptions of homotheticity and time separability. Our rejection of risk sharing based on cohort data accords with the findings on microdata of Altonji, Hayashi, and Kotlikoff (1992).

While our findings rule out the standard model of altruism in which information is symmetric, one can construct more elaborate models of altruism with information asymmetries in which clan member Euler errors do depend on the clan members' particular income realizations. For example, Kotlikoff and Razin (1988) present a model in which altruistic parents cannot observe their children's work efforts. To avoid being manipulated by children who take it easy and claim their earnings are low because they have low ability, parents will condition their transfers on the observable earnings of their children. In this model, government redistribution between parents and

children, assuming it is observable, would not change the parents' information sets and would, therefore, be completely offset by the parents in their private transfer behavior.

However, for this asymmetric information model to be plausible one would expect to see annual or at least periodic inter vivos transfers to children, with these transfers depending on the donor's income. But all the cross-section studies of U.S. private transfers suggest that, apart from the very wealthy, the vast majority of U.S. citizens neither receive nor make transfers on a routine basis. In addition, the evidence on U.S. bequests indicates that they are almost always divided equally among children with no regard to their children's current, let alone past, earnings. These facts seem strong evidence against altruism models with asymmetric information.

If one rejects, as do we, the standard as well as more elaborate varieties of the altruism model, one is left with the following implication: government redistribution between cohorts will raise the consumption of cohorts receiving the transfers and lower the consumption of those making the transfers. But what is the implication of this for national consumption and national saving? In the case of policies that redistribute from future generations to current generations, the implication is clearly an increase in national consumption and a decline in national saving. Even policies that simply redistribute between currently living cohorts will lower national saving if the redistribution runs from younger to older cohorts. The reason is that older cohorts, at least in the United States, appear to have substantially higher propensities to consume than do younger cohorts out of their remaining lifetime resources.[8]

The deleterious impact of intergenerational redistribution on national saving is more than a theoretical possibility. In the past four decades the U.S. government has engineered a massive intergenerational redistribution toward older generations of the day and away from all generations coming behind them (see Kotlikoff 1992). Simulation studies of life-cycle models suggest that such redistribution would seriously lower U.S. savings and capital formation. This is indeed what has transpired. Since 1980 the United States has been saving at about two-fifths the rate observed in the prior three decades. This reduced saving cannot be attributed to increased government consumption since government spending on goods and services relative to NNP has remained roughly constant since the mid-1950s. Rather it reflects an increase in the rate of private sector consumption out of the annual output left over after government consumption is netted out. If this increase in private sector consumption is due to redistribution toward the elderly, one would expect to observe a shift since the 1950s in the shape of the cross-section, age-consumption profile, with the relative consumption of the elderly increasing

8. This statement is based on research in progress by Jagadecsh Bokhala and Laurence J. Kotlikoff to measure how the propensity to consume varies with age.

over time. This is precisely what is emerging from preliminary analysis of the 1962–63, 1972–73, and the 1980s Consumer Expenditure Survey (see Kotlikoff, Gokhala, and Sabelhaus forthcoming).

To conclude, this essay contributes to a growing body of findings that document the effectiveness of intergenerational redistribution in raising the consumption of generations receiving income from the government and lowering the consumption of generations giving income to the government. While the precise effect of postwar U.S. intergenerational redistribution on U.S. saving may never be known, the available evidence from this and other studies suggest that past and present U.S. intergenerational redistribution may well be the main reason for the long term and continuing decline in U.S. saving.

APPENDIX: THE LIKELIHOOD FUNCTION AND THE DERIVATION OF THE ESTIMATORS

Under our assumption that the η_t^as are normal and independent across time, the log of the likelihood function, L, is given by:

$$L = \log K - \frac{1}{2}\sum_{t=1}^{T} \log|V_t| - \frac{1}{2}\sum_{t=1}^{T} \eta_t' V_t^{-1} \eta_t \qquad \text{(A.1)}$$

where $\eta_t = Y_t - \psi - \mu_t i$. The term Y_t is a column vector whose elements are $\hat{Y}_t^a$. The vector ψ captures the time-invariant, age-specific constants arising in equation 13 when $\bar{\psi}_t^a$ is time invariant. The vector i is a column vector of 1s. The term T equals the number of time periods in our data set.

The matrix V_t equals $H_t' V H_t$, and V is defiined by

$$V = \sigma^2 \begin{vmatrix} \nu^2 & \rho\nu^3 & \cdot & \cdot & \cdot \\ \rho\nu^3 & \nu^4 & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \rho^{N-2}\nu^N & \rho^{N-3}\nu^{N+1} & \cdot & \nu^{2N-2} & \rho\nu^{2N-1} \\ \rho^{N-1}\nu^{N+1} & \rho^{N-2}\nu^{N+2} & \cdot & \rho\nu^{2N-1} & \nu^{2N} \end{vmatrix}$$

where N is the number of age cohorts (58 in the case of quarterly data and 19 in the case of semiannual data) in our data, and

$$H_t = \begin{vmatrix} h_{1t} & 0 & \cdot & \cdot & 0 \\ 0 & \cdot & \cdot & & \cdot \\ \cdot & & & \cdot & 0 \\ 0 & \cdot & \cdot & 0 & h_{nt} \end{vmatrix} \quad \text{where } h_{at} \text{ equals } \sum_k w_{atk}^2 \Big/ \left(\sum_k w_{atk}^2\right) \text{ and}$$

where w_{atk} is the CES population weight of household k, which is age a at time t.

The first-order conditions resulting from maximizing equation A.1 with respect to ψ, μ_t, and σ^2 are given respectively in equations A.2, A.3, A.4:

$$\sum_{t=1}^{T} V_t^{-1} \eta_t = 0 \tag{A.2}$$

$$iV_t^{-1} \eta_t = 0 \tag{A.3}$$

$$NT = \sum_{t=1}^{T} \eta_t' V_t^{-1} \eta_t \tag{A.4}$$

From equation A.3 we have

$$\mu_t = (iV_t^{-1}i)^{-1}(Y_t - \psi) \tag{A.5}$$

Equations A.5 and A.2 imply

$$\sum_{t=1}^{T} V_t^{-1} [I - (i'V_t^{-1} i)ii'](Y_t - \psi) = 0 \tag{A.6}$$

Normalizing the sum of the μ_ts to zero yields

$$s^2 \sum_{t=1}^{T} (iV_t^{-1}i)^{-1} i' V_t^{-1}(Y_t - \psi) = 0 \tag{A.7}$$

Multiplying equation A.7 by i and adding the resulting expression to equation A.6 leads to

$$\hat{\psi} = \sum_{t=1}^{T} [V_t^{-1} - (iV_t^{-1}i)^{-1}(V_t^{-1} - I)\, i' V_t^{-1}]^{-1}$$

$$\left\{ \sum_{t=1}^{T} [V_t^{-1} - (i' V_t^{-1} i)^{-1}(V_t^{-1} - I) ii' V_t^{-1}] Y_t \right\} \tag{A.8}$$

Given knowledge of the V_ts, we can use equation A.8 plus equation A.5 to determine estimates of the μ_ts and the elements of ψ. Rather than solve analytically for the estimates of ν and ρ, we searched over a grid of alternative pairs of these parameters. For each choice of these parameters we formed the V_t matrices and used equations A.8 and A.5 to calculate the corresponding values of ψ and the μ_ts.

REFERENCES

Altonji, Joseph, Fumio Hayashi, and Laurence J. Kotlikoff. "Is the Extended Family Altruistically Linked? New Tests Based on Micro Data." *American Economic Review* (Dec. 1992): 1177–98.

Altonji, Joseph, and Aloysius Siow. "Testing the Response of Consumption to Income Changes with (Noisy) Panel Data." *Quarterly Journal of Economics* 102 (May 1987): 293–328.

Aschauer, David. "Fiscal Policy and Aggregate Demand." *American Economic Review* 75, no. 1 (March 1985): 117–27.

Auerbach, Alan J., and Laurence J. Kotlikoff. "An Examination of Empirical Tests of Social Security Savings." In *Social Policy Evaluation: An Economic Perspective*, ed. Elhanan Helpman et al. New York: Academic Press, 1983.

Barro, Robert J. "Are Government Bonds Net Wealth?" *Journal of Political Economy* 82, no. 6 (Nov./Dec. 1974): 1095–1117.

Bernheim, B. Douglas, and Kyle Bagwell. "Is Everything Neutral?" *Journal of Political Economy* 96 (1988): 308–38.

Boskin, Michael J., and Laurence J. Kotlikoff. "Public Debt and U.S. Saving: A New Test of the Neutrality Hypothesis." *Carnegie-Rochester Conference Volume Series*, February, 1986.

Boskin, Michael J., and Lawrence Lau. "An Analysis of U.S. Post War Consumption and Saving." 2 parts. National Bureau of Economic Research Working Paper, no. 2605; no. 2606. National Bureau of Economic Research, 1988.

Cochrane, John. "A Simple Test of Consumption Insurance." *Journal of Political Economy* 99 (October 1991): 957–76.

Cox, Donald. "Motives for Private Income Transfers." *Journal of Political Economy*, 95 (June 1987): 508–46.

Darby, Michael R. *Effects of Social Security on Income and the Capital Stock*. Washington, DC: Amer. Ent. Inst., 1979.

Feldstein, Martin. "Social Security, Induced Retirement, and Aggregate Capital Accumulation." *Journal of Political Economy* 82, no. 5 (Sept./Oct. 1974): 905–26.

Flavin, Marjorie A. "The Adjustment of Consumption to Changing Expectations about Future Income." *Journal of Political Economy* 89, no. 5 (October 1981): 974–1009.

Hall, Robert E. "Stochastic Implications of the Life Cycle–Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy* 86, no. 6 (December 1978): 971–87.

Hall, Robert E., and Frederic Mishkin. "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households." *Econometrica* 50, no. 2, (March 1982): 461–81.

Hayashi, Fumio. "The Effects of Liquidity Constraints on Consumption: A Cross Sectional Analysis." *Quarterly Journal of Economics*, no. 1 (Feb. 1985): 183–206.

King, Mervyn A., and Jonathan I. Leape. "Wealth and Portfolio Composition: Theory

and Evidence." National Bureau of Economic Research Working Paper, no. 1468, 1984.

Kormendi, Roger. "Government Debt, Government Spending, and Private Sector Behavior." *American Economic Review* 73, no. 4 (December 1983): 994–1010.

Kotlikoff, Laurence J. "Intergenerational Transfers and Savings." *Journal of Economic Perspectives* 2, no. 2 (Spring 1988): 41–58.

———. "Altruistic Extended Family Linkages, A Note." In *What Determines Savings?*, by Kotlikoff. Cambridge, MA: MIT Press, 1989.

———. *Generational Accounting*. New York, NY: Free Press, 1992.

Kotlikoff, Laurence J., and Assaf Razin. "Making Bequests without Spoiling Children." National Bureau of Economic Research Working Paper, no. 2735, 1988.

Kotlikoff, Laurence J., Jagadecsh Bokhala, and John Sabelhaus. "Understanding Postwar Changes in U.S. Saving." Forthcoming.

Lawrence, Emily. "Do Transfers to the Poor Reduce Savings?" Yale University, 1983. Mimeo.

Leimer, Dean R., and Selig D. Lesnoy. "Social Security and Private Saving: A Reexamination of the Time Series Evidence Using Alternative Social Security Wealth Variables." Working Paper, no. 19. Washington, DC: Social Security Administration, Office of Research and Statistics, 1980.

Mace, Barbara. "Full Insurance in the Presence of Aggregate Uncertainty." *Journal of Political Economy* 99 (October 1991): 928–56.

Mankiw, N. Gregory, Julio J. Rothenberg, and Lawrence H. Summers. "Intertemporal Substitution in Macroeconomics." *Quarterly Journal of Economics* 48: 1983.

Menchik, Paul L. "Unequal Estate Division: Is it Altruism, Reverse Bequests, or Simply Noise?" Paper presented to the 1984 Seminar on Modeling the Accumulation and Distribution of Wealth, Paris, France, 1984.

Rosenzweig, Mark R. "Risk, Implicit Contracts, and the Family in Rural Areas of Low Income Countries." *Economic Journal* 98 (December 1988): 1148–70.

Samuelson, Paul. "Lifetime Portfolio Selection by Dynamic Stochastic Programming." *Review of Economics and Statistics* 51, no. 3 (August 1969): 239–46.

Shapiro, Matthew. "The Permanent Income Hypothesis and the Real Interest Rate: Some Evidence from Panel Data." *Economics Letters* 14 (1984): 93–100.

Townsend, Robert M. "Risk and Insurance in Village India." University of Chicago, 1989. Mimeo.

Williamson, Samuel H., and Warren L. Jones. "Computing the Impact of Social Security Using the Life Cycle Consumption Function." *American Economic Review* 73, no. 5 (December 1983): 1036–52.

Zeldes, Stephen. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97 (1989).